



ELSEVIER

Contents lists available at ScienceDirect

Information Systems

journal homepage: www.elsevier.com/locate/infosys

Detection of radioactive sources in urban scenes using Bayesian Aggregation of data from mobile spectrometers



Prateek Tandon ^{a,*}, Peter Huggins ^a, Rob Maclachlan ^a, Artur Dubrawski ^a,
Karl Nelson ^b, Simon Labov ^b

^a Carnegie Mellon University Pittsburgh, PA 15213, United States

^b Lawrence Livermore National Lab Livermore, CA 94550, United States

ARTICLE INFO

Article history:

Received 18 December 2014

Received in revised form

7 October 2015

Accepted 16 October 2015

Available online 10 November 2015

Keywords:

Machine learning

Data fusion

Bayesian methods

ABSTRACT

Mobile radiation detector systems aim to help identify dangerous sources of radiation while minimizing frequency of false alarms caused by non-threatening nuisance sources prevalent in cluttered urban scenes. We develop methods for spatially aggregating evidence from multiple spectral observations to simultaneously detect and infer properties of threatening radiation sources.

Our Bayesian Aggregation (BA) framework allows sensor fusion across multiple measurements to boost detection capability of a radioactive point source, providing several key innovations previously unexplored in the literature. Our method learns the expected Signal-to-Noise Ratio (SNR) trend as a function of source exposure using Bayesian non-parametrics to enable robust detection. The method scales well in spatial search by leveraging conditional independence and locality in Bayesian updates. The framework also allows modeling of source parameters such as intensity or type to enable property characterization of detected sources. Approaches for incorporating modeling information into BA are compared and benchmarked with respect to other data fusion techniques.

© 2015 Elsevier Ltd. All rights reserved.

1. Introduction

From preventing the proliferation of nuclear weapons to assessing radiation health risks from a collapsed power plant, a nation's capability to effectively monitor radiation sources in its cities is increasingly crucial to its safety. Dirty bombs built from radioactive material or radioactive emanation from stolen medical or industrial-use isotopes are frightening but very real threat scenarios. Mobile radiation detection systems provide promise for effective nuclear search and broad area monitoring for threats in urban scenes. They aim to help law enforcement officers

detect and localize sources of radiation in complex urban environments that have frequently changing radiation landscapes.

Advances in hardware design have allowed for the ability to collect significant amounts of radiation spectrum data. One of the fundamental challenges is to automate the mining and analysis of the large amounts of sensor data that can be collected in real time to provide sensitive detection capabilities but maintain low false detection rates. Naturally occurring variability in the background radiation photon count rates as well as possible nuisance sources in an environment can cause false alarms for mobile radiation detector systems. The remedy is to account for the expectable variation in background and common potential nuisances via computational models, so a system can tell a truly threatening radioactive source of

* Corresponding author.

E-mail address: prateekt@andrew.cmu.edu (P. Tandon).

interest from a benign one. For appropriate response by authorities, a significant useful capability of algorithms is to infer properties of a detected radiation source such as its intensity or isotope type.

1.1. Related work

The problem of detecting a radioactive point source from observations of radiation spectra has been previously studied by many efforts.

A well-known early study in radiation source detection looked at the case of detecting a fully isotropic source using a mobile spectrometer [1]. Since background photon count rates are typically unknown, the study identified that characterizing and suppressing the background is an important challenge in finding the source. Interestingly, the mathematical model they used found that background and source photon counts both scale proportionally with the surface area of the sensor, making it difficult to gain from a larger detector.

To improve robustness and performance on actual collected sensor data, algorithms in the literature leveraged statistical models that could capture variability and imperfections in data in an attempt to improve detection of radioactive sources in real world settings. These formulations were based both on frequentist and Bayesian approaches and resulted in better signal separation models for signal and noise components on actual collected sensor data.

One of the popular approaches is K-Sigma, which models collected photon counts under a Poisson likelihood model that takes into account the distribution of total gross photon counts in collected radiation spectra [2]. The K specifies a detection parameter such that a spectrum is flagged as containing a source if the total counts in the spectrum is larger than K standard deviations away from the total counts in a mean background spectrum.

Building upon the anomaly detection theme, the Spectral Anomaly Detector algorithm [3] employs Principal Component Analysis (PCA) to capture major directions of background fluctuation and variation. The Spectral Anomaly Detector algorithm can be used to capture the key principal linear directions of variance in background spectra containing multiple energy bins. Projecting a new radiation observation onto a learned basis for background (and subtracting out projections in these directions) results in a spectral anomaly score, which can be used to decide whether the observation exhibits source-like behavior or is more background-like.

Both vanilla K-Sigma and the Spectral Anomaly Detector score individual sensor spectrum observations. Using Bayesian techniques, algorithms could be extended to account for fluctuations in signal and noise across multiple correlated sensor observations via Bayesian data fusion [4–6]. For instance, particle filters [7–9] are a popular approach to aggregating multiple observations to detect a target in a Bayesian framework.

Aside from the Bayesian data fusion approach, the Weighted Combining (WC) method [10] has been popular for fusing evidence from multiple observations to detect sources. The WC method uses $1/r^2$ distance weighting on

measurements to flag a source location. The method is presumed to be very powerful for flagging the locations in the environment that maximize the estimated SNR at the locations, given the source is isotropic. The algorithm maintains a geographic background “map” and a geographic source “map” which are iteratively updated. Given a new measurement, WC estimates the signal and noise components of the measurement and adds these estimates to its running estimates of signal and noise at geographic locations in its maps. The geographic location with the highest SNR score, after aggregation, is predicted to be the source location.

1.2. Innovations of Bayesian aggregation

Our Bayesian Aggregation (BA) approach builds upon many of the existing works but provides many improvements not previously explored in the literature.

First, BA provides key innovation on the radiation source/background signal separation problem. It is one of the premier Bayesian methods to fully utilize spectral information in empirical modeling of data likelihoods from real-world background data instead of just total photon counts in a spectrum. Using nonparametric density estimation techniques and appropriate measurement scoring schemes, BA can effectively suppress background radiation and non-threatening radiation emanated by nuisance sources without making *a priori* parametric assumptions about the distribution of local background and nuisances. These capabilities help provide robust signal separation when compared to other methods of aggregating evidence.

Second, our framework enables not only source localization from multiple spatially-correlated observations but inference of characteristics of the source such as the source intensity multiplier or source type. The inferential capabilities of BA allow simultaneous tracking of multiple, multi-modal hypotheses about the source parameters in a joint space while detecting the source. Modeling of source intensity and source type information in a Bayesian framework enables this previously unexplored capability.

BA resembles the framework underlying particle filtering [11,18] and sequential importance sampling [12] and, given sufficiently extensive posterior sampling, offers equivalent theoretical detection performance as particle filters [13]. However, BA is designed to be substantially more computationally scalable by using fast data structures and a manageably complex space of source parameter hypotheses. By using data structures such as kd-trees to speed up computation of hypotheses, BA can enable rapid evaluation of many real world scenarios.

2. Methods

2.1. Data

In this study, we focus on gamma-ray spectrometry measurements, where each measured photon has an associated scalar energy estimate. A single gamma-ray spectrometry measurement (after calibration) is a 128-dimensional numeric vector that histograms the energies

of received particles using 128 different energy bins, quadratically spaced across the 80–2800 keV energy range. Each measurement is collected at a particular time and location and is a 128-dimensional vector of photon counts received at the 128 monitored energy levels. Since different threats have characteristic photon energies, the use of spectra provides resolution into the type of threat being detected. The choice of 128 bins and quadratic energy binning provides maximal resolution on the threat space, putting particular emphasis on the low energy range where threats are likely to appear in the gamma-ray sensor's sensitivity band.

Our experimental set contains real field data collected over a period of five consecutive days using a double 4×16 in NaI planar spectrometer installed on a van driving in an urban area. The data contains 70,000 radiation spectrum measurements, reflective of background and any existing nuisance sources, collected at 2 s intervals while the vehicle was in motion. Annotation data recorded for each measurement include timestamp, longitude and latitude obtained from the GPS receiver, and the current speed of the vehicle.

Fig. 1 shows the GPS layout of the primary 2.5-km by 2.5-km rectangular region of collected data and partitioning of it into geographically disjoint subregions. Partitioning into subregions allowed us to use geographically disjoint sets of data for training, validating, and testing the models to avoid overfitting. Each geographic subregion was traversed multiple times during the five-day data collection effort, allowing models to capture temporal variation.

2.2. Point source simulator

A source simulator is used to inject user-supplied synthetic radiation profiles into collected background radiation data to simulate the presence of point sources. Given a geographic subregion, the simulator chooses a random location which is within a prescribed distance to at least one measurement (and not too close to surrounding measurements) and adds simulated source photon counts to the pre-existing background measurements. The simulator models injection counts, taking into account relative location of the sensor and the source, velocity of the sensor-carrying vehicle, measurement time interval, and the Poisson distribution of photon counts.

Fig. 2a shows a typical background radiation spectrum measurement, a fissile material source template, and the signal that results from the additive injection. Fig. 2b shows the effect of an example injection geographically on a test subregion of the city. An injected point source affects proximal measurements (i.e. within 20 m). Simulation can include varying the point source intensity or isotope type of injected source.

2.3. Bayesian aggregation

BA relies on field data and injected synthetic sources to learn statistical models of expected threat. Training processes are computationally affordable, and models can be cheaply and frequently re-trained if needed (for instance,

to adjust to new environment backgrounds). Trained models are used to process sequences of new spectral measurements taken along the route of the mobile sensor. The BA pipeline has three key stages as shown in Fig. 3.

BA receives as input radiation spectra and map locations of measurements. The first stage of BA is to estimate the Signal-to-Noise Ratio (SNR) of a measurement in terms of its source signal and background noise components. Once the SNR is estimated, location, velocity, and other positioning information can help quantify the expected exposure to a source. The second stage of BA builds a probabilistic sensor model that can score whether the measurement follows the expected exposure-SNR trend for a point source at hypothetical source locations on the map. Finally, evidence is spatially aggregated across multiple observations using Bayesian data fusion to robustly test these hypothetical source locations and render a threat probability map.

2.3.1. Training measurement SNR estimators

The first phase of the Bayesian Aggregation algorithm involves training a SNR estimator for single measurements using the geographic region allocated for “training” in our field data. A radiation measurement M_p (a 128×1 vector), collected at a particular location and time, is presumed to be composed of photon counts due to background and source components. The p subscript on the measurement indicates a different time and geographic location for each single spectral measurement. The scalar total sum of photon counts in a collected spectrum observation, Y_p , is presumed to be given by

$$Y_p = B_p + S_p \quad (1)$$

where B_p is the scalar number of total photon counts from background in that spectrum and S_p is the scalar number of total photon counts from the source of interest in that spectrum. Note that S_p can be zero if the particular spectral observation M_p is unaffected by the source.

The goal of SNR estimation is to estimate the quantities B_p and S_p in M_p . The scalar Signal-to-Noise Ratio (SNR) of




Fig. 1. Division of the available urban data into 8 geographically disjoint subsets. These subsets are chosen as training, validation, and test disjoint subsets of data to avoid overfitting while performing experiments.

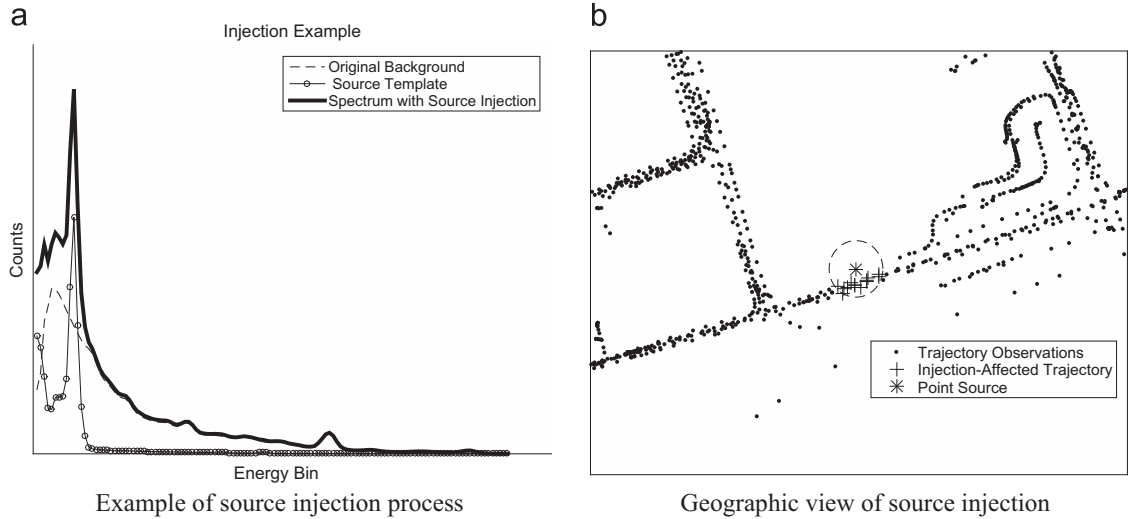


Fig. 2. Snapshot of radiation simulator capabilities. (a) shows an example of the source injection process into background radiation, while (b) provides a geographic view of source injection in an example geographic subregion of the data.

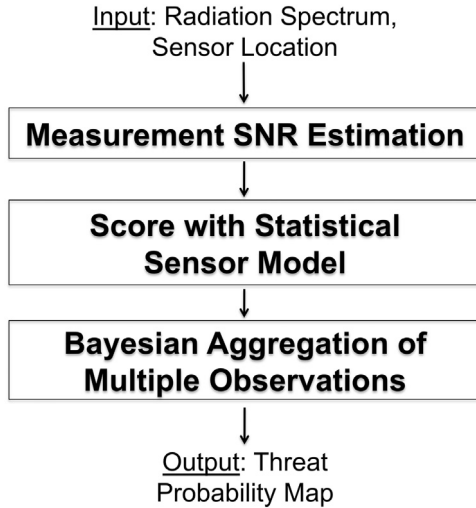


Fig. 3. Stages of BA pipeline.

the measurement M_p can then be defined as:

$$SNR(M_p) = \frac{S_p}{\sqrt{B_p}} \quad (2)$$

The signal term is given by the source signal estimate, and the noise term is formed from the background estimate. Due to the Poisson character of the data generating process, $\sqrt{B_p}$ is the noise term in the SNR [14].

We experimented with two types of measurement SNR estimators commonly used in practice: anomaly detectors and match filters. The goal of anomaly detection is to flag spectra that are distinct from typical background. This approach assumes no knowledge of a source template and simply uses a model of expected background variability to measure the background (typical variation) and source (unusual deviation) components of a measurement. Match filtering, in contrast, allows the user to specify a source template or design to match spectrum observations

against. Both approaches estimate S_p and B_p to score the SNR of a measurement.

In our anomaly detection approach, we leverage the well-known technique of Principal Components Analysis (PCA) to learn a model for typical background [15]. If we stack the P spectral measurements $M_{p \in P}$ into a $P \times 128$ data matrix, X , PCA calculates the eigendecomposition of $X^T X$ using Singular Value Decomposition:

$$X^T X = V \Sigma V^T \quad (3)$$

where V contains the eigenvectors of the original data. The top few eigenvectors (or “principal components”) of V are the linear directions that retain most of the background variation of the original data. These identified components of background can subsequently be removed from new spectral observations, leaving an estimate of source signal.

In our experiments, we use the Spectral Anomaly Detector version of PCA [3] that estimates V using the correlation matrix of background data rather than the covariance. This ensures that all energy features are on the same scale and avoids overly biasing the PCA model towards the most common background energies. The resulting compressed space is used as the null-space model, capturing the expected types of variability in the background data. We retain the top five principal components (PCs) in V since these have been empirically shown to capture the principal modes of variation in the radiation data, with more components not making a significant difference [3].

The 128×5 matrix of the top 5 PCs in V , $V_{1:5}$, can be used to estimate the SNR of new measurements. The data projection of a measured spectrum M_p (where M_p is a 128×1 vector) on the learned PCA basis is given by $T = M_p V_{1:5}$. The total sum of squared reconstruction error over all bins serves as the scalar source estimate, $S_p = \|M_p - TV_{1:5}\|_2$. The spectral anomaly detector treats the noise estimate, $\sqrt{B_p}$, as proportional to the total counts in the spectrum and simply uses it as a normalizer. $\sqrt{B_p} = \|M_p\|_1$. The L_2 norm is used for the source estimate

to maximally highlight anomalies, while the L_1 norm is used for the noise estimate simply to provide a normalizer in terms of the background for SNR. The SNR score for the measurement is $SNR(M_p) = S_p / \sqrt{B_p}$.

Match filtering [10] is an alternate way to estimate the SNR of a measurement. Match filtering allows the user to specify a source template in estimating source amount. Based on the source template, a particular range of energies (called the source window W) is monitored as likely containing source signal. Match filtering estimates source amount by linearly regressing the background photon amounts outside the source window (call those energy bins \bar{W}) to predict the number of background photon counts inside the window (W). Removing the background from the source window leaves mostly source signal there.

Training a match filter involves the following steps:

1. The user splits the energy space into two subsets: W , a feature-selected source window where characteristic source peaks (in the source template) are likely to appear and \bar{W} , the complementary set of energy bins that are not in the source window (and presumed to contain no source counts). Note that these subsets of the energy space may be any arbitrary disjoint subsets such that $W \cup \bar{W}$ forms the full 128-dimensional energy space.
2. A regression estimator is trained to predict, for a new radiation vector, the background content in the source window (W) from background counts outside the source window (in the \bar{W} energy bins). A data set of training spectra with known background amounts is assembled. The standard Least Squares Estimator is fit to the data:

$$\hat{B}(y) = (X^T X)^{-1} X^T y \quad (4)$$

X is the $N \times |\bar{W}|$ data matrix of predictor energy bins outside the source window. Each row of the X matrix is a training data spectrum that keeps only energy bins in \bar{W} . y is the $N \times 1$ vector of total sums of photon counts in the source window for each of the N training spectra. The regression estimator predicts the number of background counts in the source window (y) from the vector of photon energies outside the window (X).

Given a new observed spectrum vector, M_p , define M_p^W as the $|W| \times 1$ vector of photon counts in source window energy bins in M_p and $M_p^{\bar{W}}$ as the $|\bar{W}| \times 1$ vector of photon counts of energy bins outside the source window in M_p . Note that $\text{length}(M_p^W) + \text{length}(M_p^{\bar{W}}) = \text{length}(M_p) = 128$.

The scalar background estimate is given by $B_p = \hat{B}^T M_p^{\bar{W}}$, the regression prediction of the background amount in the source window from the background in energy bins outside the source window. The scalar source estimate is given by $S_p = |M_p^W|_1 - B_p$ where $|M_p^W|_1$ is the scalar sum of total photon counts in the source window. The SNR score is formed as $SNR(M_p) = S_p / \sqrt{B_p}$.

2.3.2. Learning probabilistic sensor models

After training a measurement SNR estimator (either an anomaly detector or a match filter), the second phase of BA

training involves building probabilistic sensor models of expected SNR score distributions as a function of source exposure. The probabilistic sensor models are estimated from actual (and source-injected) field data. In order to avoid overfitting, BA uses the “validation” subregion of field data to assemble distributions of SNR scores for positive observations (source-injected background data) and negative observations (pure background data). The score distribution for negative data forms the *null* distribution (H_0), and the score distribution obtained for source-injected data becomes the *alternate* (H_1) probability distribution.

Each distribution is parameterized by the total source exposure statistic computed from vehicle velocity, hypothesized source intensity, and relative locations of the source and the vehicle. The source exposure statistic is a multiplier for the Poisson parameters of injected source counts. For a detector of negligible volume approximated as a point, the source exposure statistic can be defined as proportional to $\int dt/R^2$, where t is time, R is the distance between source and detector, and the integral is computed over the duration of the measurement as the detector moves along its trajectory. The source exposure statistic can, in principle, be computed taking known or anticipated occlusions into account.

The two-dimensional probability densities of expected SNR as a function of the exposure statistic can be robustly estimated using nonparametric density estimation. The two-dimensional Kernel Density Estimator (KDE) [16] is often used to perform the density estimation:

$$\hat{f}_h(x_1, x_2) = \frac{1}{N} \sum_{i=1}^N \frac{1}{h_1 h_2} K\left(\frac{x_1 - X_{i,1}}{h_1}, \frac{x_2 - X_{i,2}}{h_2}\right) \quad (5)$$

where (x_1, x_2) is a query point of the exposure/SNR estimated density, $K(u_1, u_2)$ is a kernel function, X is the $N \times 2$ matrix of validation data point exposure statistics and SNR estimates, and h_1 and h_2 are bandwidth parameters that control smoothing amount. The dimensions x_1 and x_2 of our models are the exposure and SNR respectively. We use the boxcar kernel and bandwidth choices of 1. Choices of other standard kernels and bandwidth values did not significantly influence results.

The SNRs and total exposure values for positive observations are computed for each observation in validation data, and density estimation is applied to obtain a two-dimensional matrix of probability density values, with axes given by SNR and the exposure statistic. For each value of the exposure statistic, the corresponding row of the matrix gives the estimated SNR distribution. The density matrix is the basis for our estimate of $P(D|H_1)$, the probability of observing a set of given measurements D under the alternate hypothesis that a source of a particular intensity and type is present at a particular location. Similarly, a density matrix for negative observations produces an estimate of $P(D|H_0)$, the probability of observing given measurements under the null hypothesis that there is no source at that location.

Fig. 4 shows example null and alternate hypothesis models obtained for field data for representative anomaly detector and match filter SNR estimates. For null models,

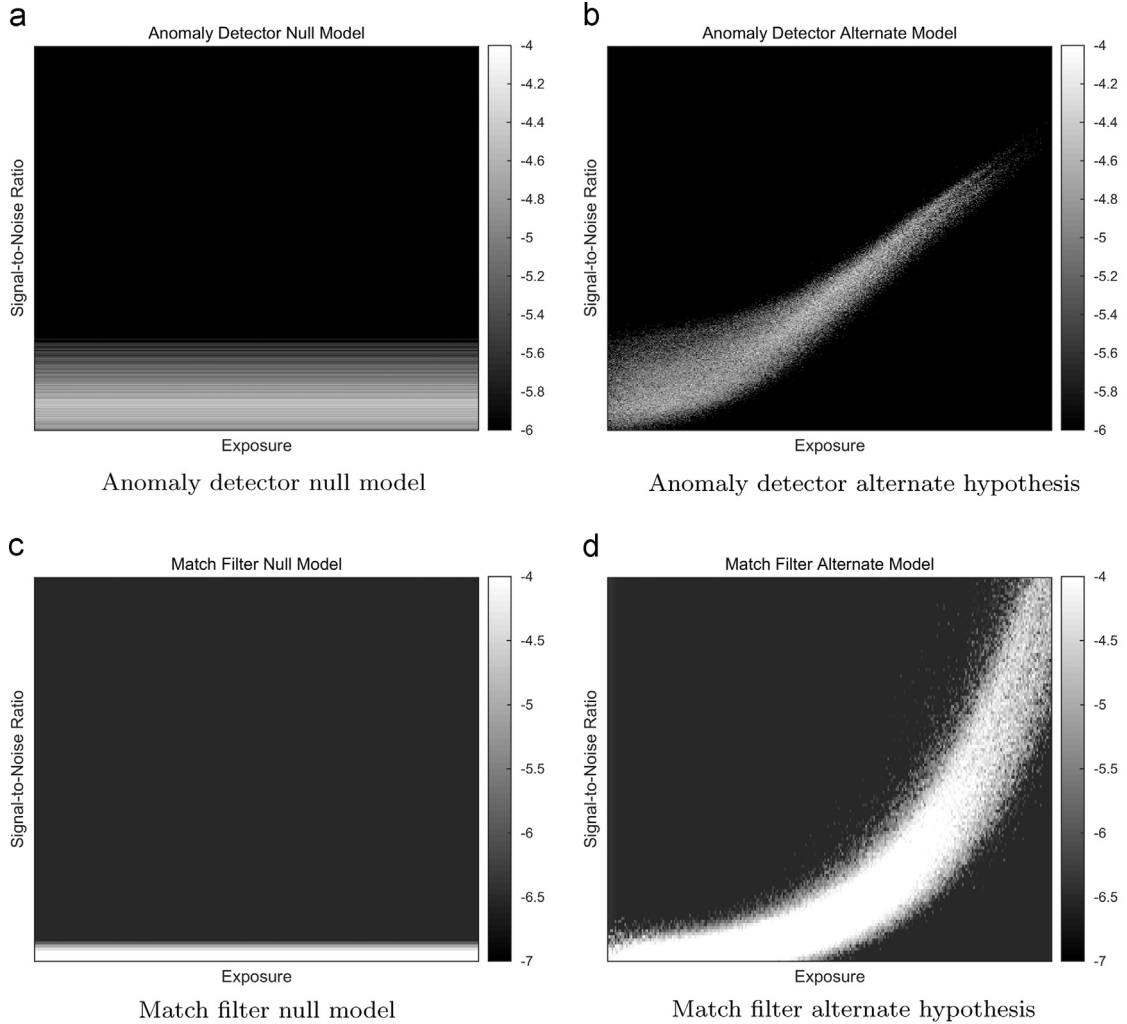


Fig. 4. Empirical sensor models for null and alternative hypothesis based on exposure and SNR. The difference in probability mass between the null and alternate densities allows for sensitive source detection. (a) Anomaly detector null model. (b) Anomaly detector alternate hypothesis. (c) Match filter null model. (d) Match filter alternate hypothesis.

SNR is flat as a function of exposure since there are no sources (of the threatening type injected into the data) to be exposed to. For alternate models, as exposure to a source goes up, so does the expected SNR. An advantage of BA is that it can capture error and variation of estimators in models nonparametrically. There is potential structure in these distributions (such as intricate double tail behavior) that would be hard to capture well by a parametric model of manageable complexity.

2.3.3. Spatial aggregation of evidence

The next step of BA is to spatially combine evidence as it is collected. Alternate source hypotheses H_1 state that a particular location in the environment contains a source of a particular intensity and type, and null hypotheses H_0 state that no source with those parameters is present. For a given terrain, the scene can be covered with a set of hypothetical source locations (e.g. distributed over a regular planar grid). As new measurements are collected and added to the overall data D , BA maintains and updates

estimates of the probabilities $P(H_1|D)$ for each source hypothesis and each null hypothesis $P(H_0|D)$.

The exposure and SNR are estimated for individual measurements D_j . These serve to index into the KDE function and look up the probabilities $P(D_j|H_0)$ and $P(D_j|H_1)$. For a particular hypothesis \mathbf{H} (either a H_1 or H_0), BA is motivated by the following equation for posterior probabilities $P(\mathbf{H}|D)$, assuming conditional independence of measurements:

$$P(\mathbf{H}|D) \propto P(\mathbf{H}) \prod_{D_j \in D} P(D_j|\mathbf{H}) \quad (6)$$

where $P(\mathbf{H})$ is the prior probability (belief) assigned to \mathbf{H} . For sufficiently low values of source exposure statistic, $P(D_j|H_1) \simeq P(D_j|H_0)$. This motivates the following algebraic manipulation when evaluating the Likelihood Ratio between H_1 and H_0 :

$$\frac{P(H_1|D)}{P(H_0|D)} \simeq \frac{P(H_1)}{P(H_0)} \prod_{D_i \in D} \frac{P(D_i|H_1)}{P(D_i|H_0)} \simeq \frac{P(H_1)}{P(H_0)} \prod_{D_k \in D \wedge D_k \in D'} \prod_{D_l \in D \wedge D_l \in D'}$$

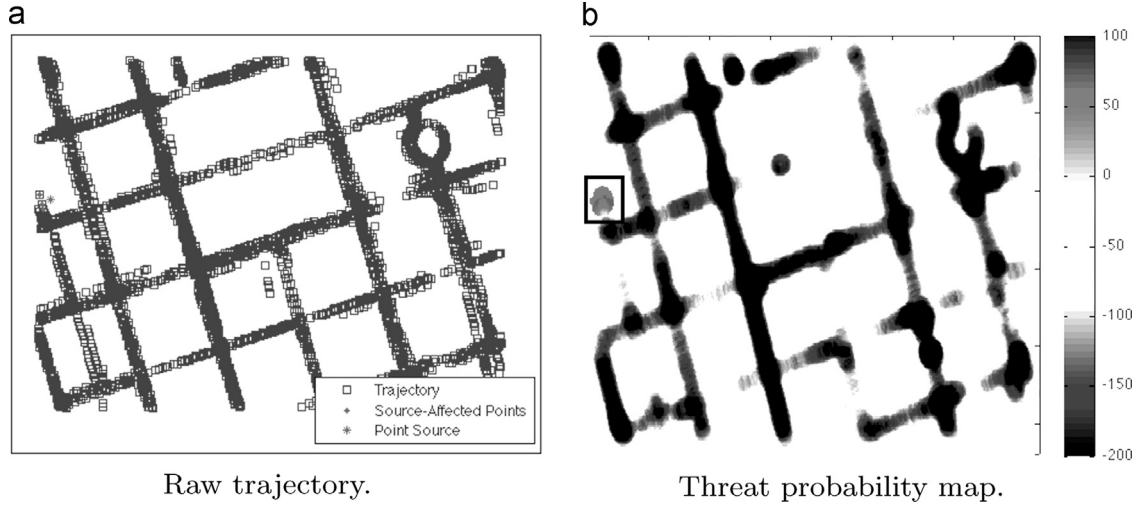


Fig. 5. Example data processing through BA. (a) shows the raw trajectory with a point source affecting measurements. (b) shows resulting threat probability map in BA. The boxed region indicates a correctly detected radiation source.

$$\frac{P(D_k|H_1)}{P(D_k|H_0)} \prod_{D_j \in D \wedge D_j \neq D} \frac{P(D_j|H_1)}{P(D_j|H_0)} \approx \frac{P(H_1)}{P(H_0)} \prod_{D_k \in D'} \frac{P(D_k|H_1)}{P(D_k|H_0)} \quad (7)$$

where $D' \subset D$ is the subset of measurements whose value of source exposure statistic is sufficiently high for a source location. The algebraic manipulation shows that, for a new measurement, only hypothetical source locations within a particular distance of it (e.g. 20 m) need to have their posterior scores updated as the measurements with low exposures cancel from the likelihood ratio. For each new measurement, kd-trees [17] efficiently find the set of hypothesized source locations locally affected by the measurement. This enables more scalable computation of Bayesian probabilities than cited previous Bayesian methods.

Fig. 5 shows an example result of BA. Fig. 5a shows physical locations where measurements were made and the location of an injected point source. Fig. 5b shows the mapping of threat probabilities by BA. It shows correct detection and localization of the injected synthetic source.

3. Experimental setup

Using the point source simulator described earlier, we repeatedly injected an isotropic point source at a random location in the subregion off the road with a fixed intensity. Source locations were restricted to be at least 10 m away from the centers of the roads to simulate realistic road-side source conditions.

We benchmarked different versions of our BA algorithm against each other as well as in comparison to an alternative method of evidence aggregation currently used in the field called the Weighted Combining (WC) Method [10]. The WC method scores source location hypotheses by estimating the SNR at each location, aggregating evidence

using the following update rules:

$$Map_S(x, y) = Map_S(x, y) + \frac{S_p}{R(x, y)^2}$$

$$Map_B(x, y) = Map_B(x, y) + \frac{B_p}{R(x, y)^4}$$

$$SNR(x, y) = \frac{Map_S(x, y)}{\sqrt{Map_B(x, y)}} \quad (8)$$

WC maintains a geographic “source map” ($Map_S(x, y)$) and a geographic “background map” ($Map_B(x, y)$) containing hypothetical source locations (x, y) . For a new measurement M_p , SNR estimators are used to compute S_p and B_p for that measurement. These quantities are weighted by the measurement’s $1/r^2$ exposure statistic and added to running totals. WC is a fierce competitor method, as it is designed to weight and aggregate measurements to maximize expected SNR in predicting whether an isotropic point source exists at location (x, y) or not.

We compared our Bayesian method of aggregation to the WC method using both the anomaly detection and match filter estimators. The metrics of success used were (1) detection performance, (2) localization accuracy, and (3) accuracy of source intensity and source type inference.

In order to facilitate the evaluation, a grid of hypothetical source locations was overlaid at 2 m resolution within 20 m of the set of field measurements. Each algorithm had to score each grid location as likely to contain a source. The grid of hypothetical source locations was scored with and without the presence of an injected point source.

A true positive is defined as the top scoring grid point within 40 m of the true injected source when injection of source counts is performed. A false positive is defined as the top scoring grid point in the same 40 m disc around the source location when injection is not performed. Detection success is measured using Receiver Operating Characteristic (ROC) curves using true positives and false positives collected over a set of source location simulations. The 40 m distance was used to avoid symmetry

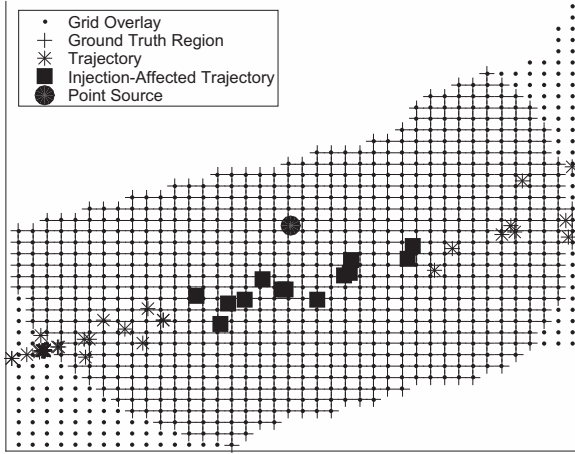


Fig. 6. Setup for experimental evaluation.

effects in localizing a source. When using an omnidirectional sensor on a perfectly straight road, the trajectory does not give information about which side of the road the source may be on. Thus, if the source is injected 10–20 m off the road, the top scoring hypothesis for the source is degenerated to a maximum of 40 m around the trajectory. Fig. 6 shows an example of our evaluation setup.

Localization accuracy is quantified by the distance of the top-scoring source location hypothesis to the true source location and summarized with a distance cumulative distribution function (CDF).

Accuracy of inference of source parameters is summarized in confusion matrices parameterized with a particular setting of the false positive rate (FPR). In our experiments, the numbers reported in confusion matrices reflect the fraction of correct inferences over the total number of experiments at FPR=0.01%. This is a commonly used metric since the beginning of the ROC curve (FPR=0.00%) is prone to noise and less stable than the FPR=0.01% mark. In our set of at least 1000 false positives, FPR=0.01% was always a stable cutoff. This cutoff is equivalent to the average interval between consecutive false alerts being 2 h 46 min 40 s of system operation.

4. Results and discussion

10,000 instances of the same test area of the city were prepared, each with a single injected point source of a fixed intensity. Fig. 7 plots the ROC detection curve comparing BA and WC with anomaly detector and match filter SNR estimators. In our experiments, match filtering outperforms anomaly detection with regularity simply because it is provided with the exact knowledge of the source type template.

BA outperforms WC regardless of choice of SNR estimator. BA captures variability in SNR as a function of expected exposure nonparametrically. It can outperform WC, which makes a Gaussian assumption about the SNR estimates. WC cannot capture the more intricate structures in their distribution. Furthermore, the WC method is affected by occasional nuisance sources in the data that

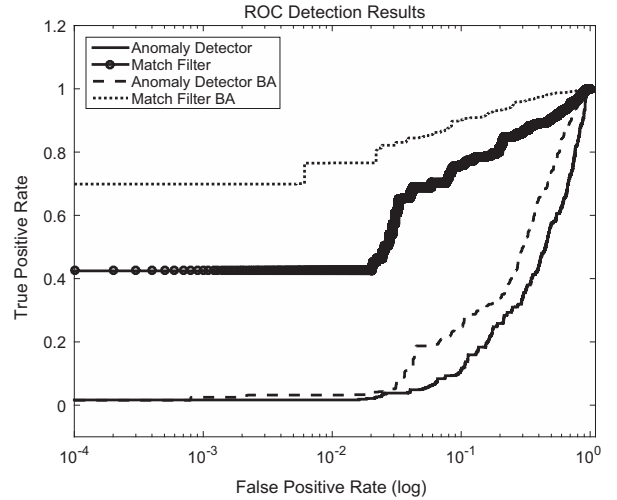


Fig. 7. Receiver Operating Characteristic (ROC) curve results for algorithms.

lead to erroneous SNR estimates. The exposure-SNR sensor models in BA help discount such noise.

We compared several methods of including source intensity and source type information in BA. The simplest method is BA-Specific which trains a separate model for each discrete setting of a parameter. Thus if there are four expected source intensities, $\{I_1, \dots, I_4\}$, there would be four distinct alternate hypothesis BA models $\{M_{I_1}, \dots, M_{I_4}\}$, each trained on data for one specific intensity parameter setting. BA-Specific methods assume $P(H_1|D) = P(H_1|D, M_{I_j})$. BA-Agnostic trains a single model learned from data containing all possible settings of the parameter. Thus, for our four intensity example, there would be only a single BA-Agnostic model, M_{I_1, \dots, I_4} , learned from a data set that mixes simulated data from all four intensities and is effectively “agnostic” to any setting of the parameter. BA-Agnostic assumes $P(H_1|D) = P(H_1|D, M_{I_1, \dots, I_4})$.

Posterior distribution information from an ensemble of BA-Specific models can also be used to detect a source. BA-Max maximizes over the posterior intensity hypothesis distribution to score the alternate hypothesis:

$$P(H_1|D) = \max_j P(H_1|D, M_{I_j}) \quad (9)$$

BA-Marg, in contrast, marginalizes over the posterior intensity hypothesis distribution to score the alternate hypothesis:

$$P(H_1|D) = \sum_j P(H_1|D, M_{I_j}) \quad (10)$$

Four different settings of the intensity multiplier were used, increasing in a geometric series, to simulate increasingly pronounced sources. Sources are well in the tolerance of background. At the lowest intensity setting, source stand-off count rate at the distance of 10 m is 150 photons per second, while mean background count rate is 1263 per second with standard deviation of 267 photon counts per second. In the energy bins selected by the match filter as being in the source window, the stand-off

count rate at 10 m for the lowest intensity is 31 per second, and the mean background count rate is 125 per second with a standard deviation of 31 per second. All conversions to source count rate for the four used intensity multipliers are shown in Table 1.

We created four separate sets of 1000 simulated sources for the same test subregion of the city, one set of simulated injections for each setting of source intensity. Then we compared performance of algorithms on these simulated worlds. Fig. 8a–d shows ROC detection results for the different algorithms for the four settings of the true source intensity.

Table 1
Intensity multiplier to count rate conversion.

Intensity	Source count rate	Feature-selected source count rate
I1	150 counts/s	31 counts/s
I2	196 counts/s	77 counts/s
I3	255 counts/s	101 counts/s
I4	331 counts/s	131 counts/s

The results indicate that all BA methods (using match filter) outperform the WC method (that also uses match filter). Since the WC method does not specifically model source intensity, the BA methods trained with the correct intensity model outperform it easily. The WC method is prone to flagging nuisance sources or be fooled by a weak source. In contrast, intensity-aware BA can provide reliable detection of sources of unknown intensity.

A similar experiment was performed with the source type parameter. Three different source templates were chosen for injection. Fig. 9 shows example injections of the source templates into typical background radiation.

1000 instances of the same test block were simulated with different injected source types. Fig. 10a–c shows ROC trends in detecting the three different source types.

Once again, BA methods do well against the WC method with match filtering. BA-Marg and BA-Max have little loss when compared to the BA-Specific model trained for the true source type. BA-Marg is the method of choice since marginalization over the intensity and source type variations provides not only the theoretically optimal detection power but promising empirical results as well.

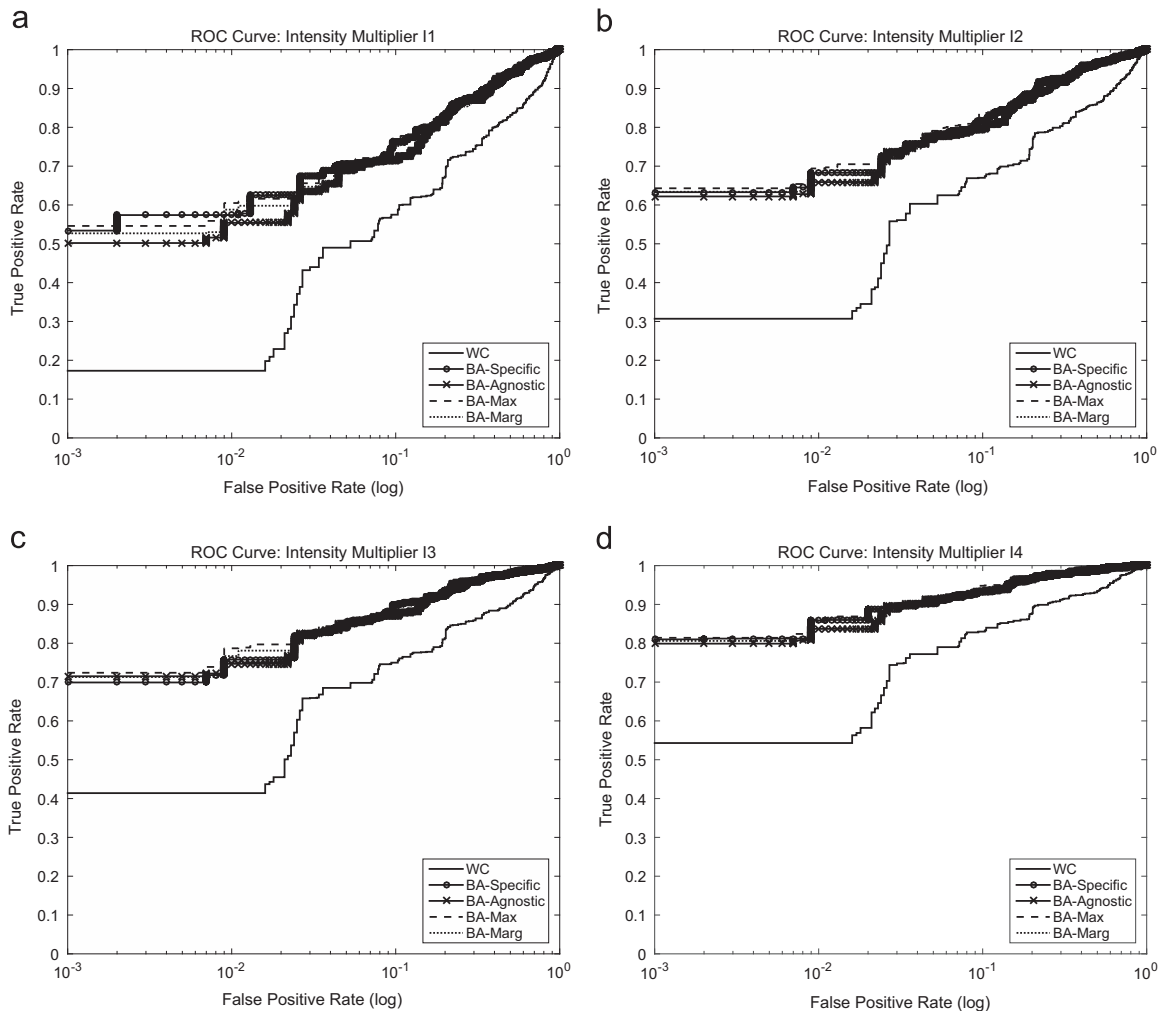


Fig. 8. ROC diagrams for source intensity experiments. All BA methods outperform WC.

Fig. 11 shows the CDF of source localization error for performance of BA-Marg in the presence of intensity variations (Fig. 11a) and source type variations (Fig. 11b). The

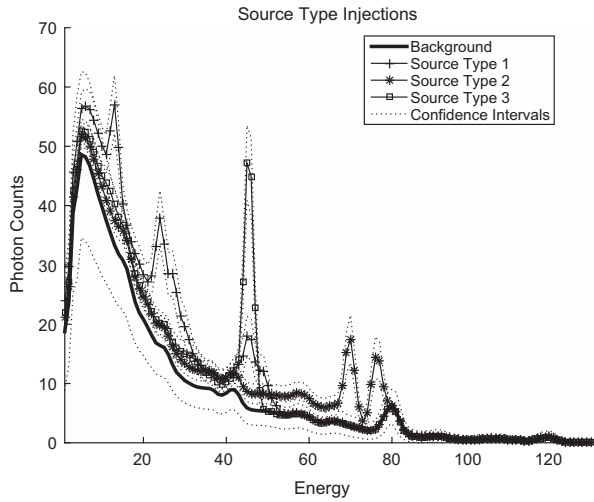


Fig. 9. Source templates injected into background.

algorithm can robustly localize many sources of varying intensity and source type, even in cases where WC may fail to detect them at all.

In addition to helping provide robust detection and localization of a radioactive point source, the posterior probabilities from the BA models can be used to infer the source intensity and source type parameters for the source. To evaluate accuracy of such inference capability, we compared the parameter inferred at the hypothesized source location at 0.01% false positive rate for each of the evaluated synthetic environments.

Table 2 shows intensity inference results for our two methods of inferring a parameter – maximizing over the BA-specific posterior hypothesis distribution (BA-Max) and marginalizing over the posterior hypothesis distribution (BA-marg). BA-Max slightly outperforms BA-Marg in intensity inference, though differences are not significant. Intensity is a quite weakly constrained parameter since there is a fair bit of location leeway in detecting a source within a 40 m radius. A detection algorithm can model a source as either being a strong source far away or a weak source close by and still succeed in detecting it, though yield incorrect identification of source intensity.

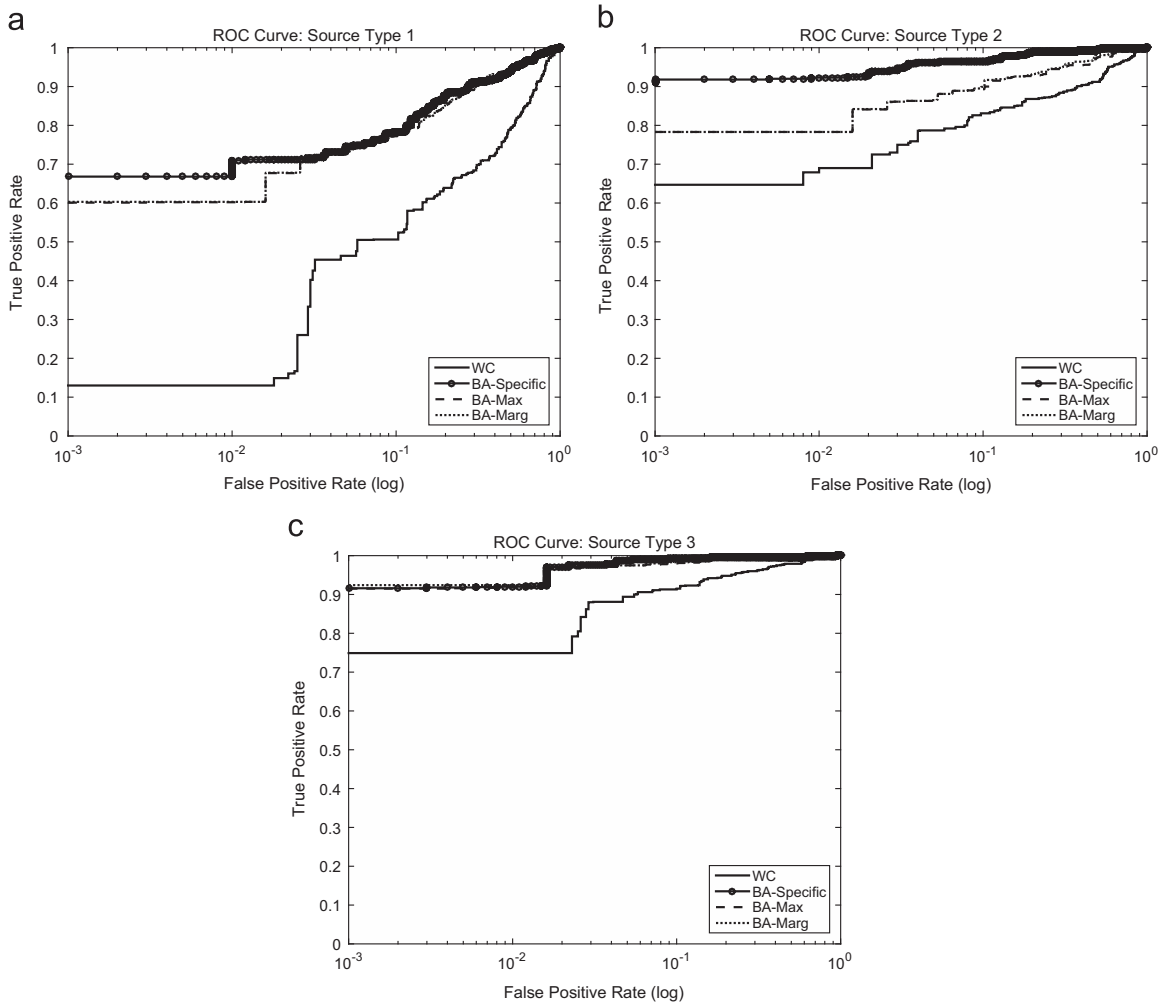


Fig. 10. ROC diagrams show threat detection results for source type experiments.

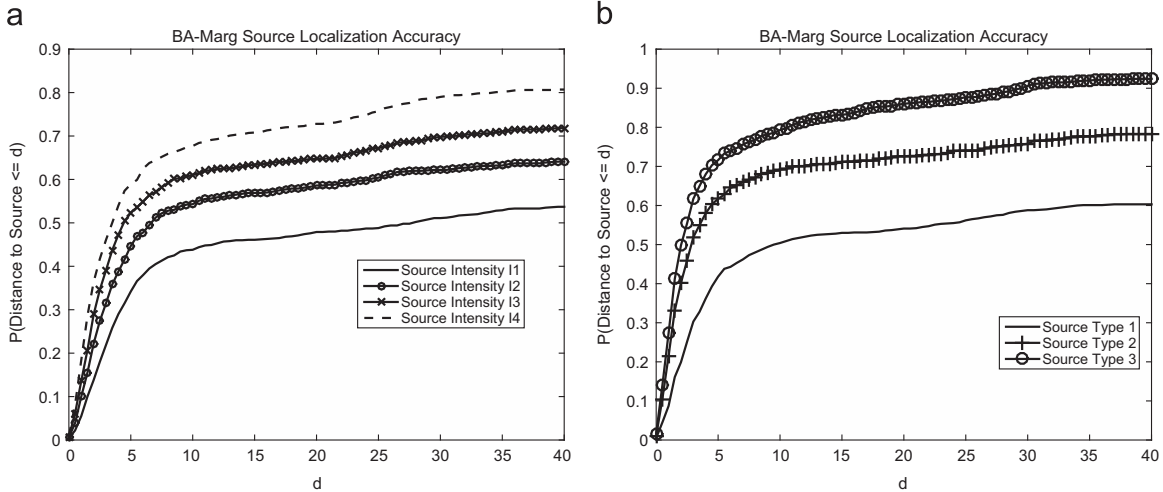


Fig. 11. Source localization accuracy results summarized with the distance from the true source location CDFs for intensity variations (a) and source type variations (b).

Table 2
Confusion matrix for intensity inference: BA-Marg/BA-Max.

True intensity →	I1	I2	I3	I4
Correct inference	35.4%/36.2%	20.7%/21.2%	21.4%/21.5%	46.8%/47.0%
Incorrect inference	23.4%/24.3%	47.4%/48.2%	55.3%/57.2%	38.7%/9.1%
False detection first	41.2%/39.5%	31.9%/30.6%	23.3%/21.3%	14.5%/13.9%

Table 3
Confusion matrix for source type inference: BA-Marg/BA-Max.

True source Type →	Source 1	Source 2	Source 3
Correct inference	57.8%/57.6%	77.5%/77.5%	87.0%/86.3%
Incorrect inference	2.50%/2.50%	0.80%/0.80%	5.40%/5.20%
False detection first	39.7%/39.9%	21.7%/21.7%	7.60%/8.50%

Table 3 shows source type inference results for BA-max and BA-marg. Results are not significantly different between algorithms, though BA-Marg slightly outperforms BA-max in the experiment. Both algorithms produce accuracies within the 58–87% range, showing BA’s capability to tell different source types apart robustly.

5. Conclusion

Our results indicate that Bayesian Aggregation of data can improve the detection of radioactive sources as well as help determine their characteristics such as location, intensity, and source type.

With regards to detection, our study illustrates how BA can boost the performance of anomaly detector and match

filter estimators commonly used among the radiation sensing community. BA captures the distribution of expected SNR as a function of exposure without making parametric assumptions about expected distribution. This approach allows for more effective dismissal of background nuisance sources that might otherwise thwart the capability of such SNR estimators in practice.

In terms of scalability, BA leverages an algebraic manipulation and conditional independences in its Bayesian updates to enable rapid calculation of posterior probabilities. The use of data structures such as kd-trees helps speed up computation of the threat probability map. The speed improvement enables efficient maintenance of many, multi-modal hypotheses about source parameters.

With regards to determining properties of detected sources, we have shown how intensity and source type information can be incorporated into BA models. Our experimentation with various methods suggests that our method of choice, BA-Marg, can simultaneously provide robust detection of sources of different intensities and source types (as evaluated by ROC curves) while inferring their characteristics (as evaluated by accuracy on confusion matrices).

BA’s algorithmic enhancement of detection capability enables mobile radiation detection systems to provide more sensitive and precise nuclear threat detection in urban scenes. Systems will be able to more accurately classify threats from non-threats, while lowering the false alarm rates. BA’s source parameter inference capabilities allow law enforcement officers to have real-time knowledge of the properties of threatening (or non-threatening) radiation sources discovered in the field, allowing for appropriate action.

Both the detection and knowledge-providing capabilities of BA seek to empower law enforcement officers in having a prompt, suitable, and effective response in dealing with radiological threats.

Acknowledgments

This work has been partially supported by the United States Department of Homeland Security, Domestic Nuclear Detection Office, under competitively awarded Grant 2010-DN-077-ARI040-02; the United States Department of Defense, Defense Threat Reduction Agency, under Award HDTRA1-13-1-0026; the United States Department of Energy, under Grant DE-NA0001736; and the National Science Foundation under Awards 0911032 and 1320347. Lawrence Livermore National Laboratory is operated by Lawrence Livermore National Security, LLC, for the United States Department of Energy, National Nuclear Security Administration under Contract DE-AC52-07NA27344. This support does not constitute an express or implied endorsement on the part of the United States Government.

Appendix A. Supplementary data

Supplementary data associated with this paper can be found in the online version at <http://dx.doi.org/10.1016/j.is.2015.10.006>.

References

- [1] K. Zioc, W. Goldstein, The lost source, varying backgrounds and why bigger may not be better, in: *Unattended Radiation Sensor Systems for Remote Applications*, vol. 632, AIP Publishing, Washington, DC (USA), 2002, pp. 60–70.
- [2] K.M. Chandy, J. Bunn, A. Liu, Models and algorithms for radiation detection, in: *Modeling and Simulation Workshop for Homeland Security*, 2010, pp. 1–6.
- [3] K. Nelson, S. Labov, Detection and Alarming with Sords Unimaged Data: Background Data Analysis, Lawrence Livermore National Lab Technical Report.
- [4] M. Morelande, B. Ristic, A. Gunatilaka, Detection and parameter estimation of multiple radioactive sources, in: *2007 10th International Conference on Information Fusion*, 2007, pp. 1–7. <http://dx.doi.org/10.1109/ICIF.2007.4408094>.
- [5] M. Morelande, B. Ristic, Radiological source detection and localisation using Bayesian techniques, *IEEE Trans. Signal Process.* 57 (11) (2009) 4220–4231, <http://dx.doi.org/10.1109/TSP.2009.2026618>.
- [6] A. Liu, J. Bunn, K. Chandy, Sensor networks for the detection and tracking of radiation and other threats in cities, in: *2011 10th International Conference on Information Processing in Sensor Networks (IPSN)*, 2011, pp. 1–12.
- [7] Z. Chen, Bayesian filtering: from Kalman filters to particle filters, and beyond, *Statistics* 182 (1) (2003) 1–69.
- [8] B. Ristic, M. Morelande, A. Gunatilaka, A controlled search for radioactive point sources, in: *2008 11th International Conference on Information Fusion*, IEEE, Cologne, German, 2008, pp. 1–5.
- [9] B. Ristic, A. Gunatilaka, Information driven localisation of a radiological point source, *Inf. Fusion* 9 (2) (2008) 317–326. <http://dx.doi.org/10.1016/j.inffus.2007.06.004>. URL <<http://www.sciencedirect.com/science/article/pii/S1566253507000577>>.
- [10] K. Nelson, S. Labov, Aggregation of Mobile Radiation Data, Lawrence Livermore National Lab Technical Report.
- [11] A. Doucet, A. M. Johansen, A tutorial on particle filtering and smoothing: Fifteen years later, in: *Handbook of Nonlinear Filtering*, vol. 12, 2009, pp. 656–704.
- [12] K.N. Borozdin, A.V. Klimenko, W.C. Priedhorsky, N. Hengartner, C.C. Alexander, R.A. Cortez, H.G. Tanner, X. Papageorgiou, Optimized strategies for smart nuclear search, in: *IEEE Nuclear Science Symposium Conference Record*, 2006, vol. 2, IEEE, San Diego, CA, 2006, pp. 926–928.
- [13] J.K. Kruschke, *Doing Bayesian Data Analysis: A Tutorial with R and BUGS*, Academic Press, Burlington, MA, 2011 URL <http://www.worldcat.org/search?qt=worldcat_org_all&q=0123814855>.
- [14] G.F. Knoll, *Radiation Detection and Measurement*, John Wiley & Sons, 2010.
- [15] C.M. Bishop, et al., *Pattern Recognition and Machine Learning*, vol. 1, Springer, New York, New York (USA), 2006.
- [16] T. Hastie, R. Tibshirani, J. Friedman, T. Hastie, J. Friedman, R. Tibshirani, *The Elements of Statistical Learning*, vol. 2, Springer, Berlin, Germany, 2009.
- [17] D.P. Mehta, S. Sahni, *Handbook Of Data Structures And Applications (Chapman & Hall/Crc Computer and Information Science Series.)*, Chapman & Hall/CRC, Boca Raton, Florida (USA), 2004.
- [18] T.B. Schön, F. Gustafsson, R. Karlsson, *The Particle Filter in Practice*, 2009.